Latency Amplification Through Token Quest for DDoS Mitigation

25 June 2025

Simon Edwards
Research Acceleration Initiative

Introduction

Protecting against Distributed Denial of Service attacks has become a multi-billion dollar industry, but even the most sophisticated DDoS mitigation services cannot effectively protect against DDoS attacks by nation-state actors. Novel techniques continue to be coveted for mitigating the effectiveness of DDoS attacks.

Abstract

A novel approach for mitigating the effectiveness of DDoS attacks is something which may be termed Latency Amplification through Token Quest (LATQ.) A system utilizing LATQ would, upon receiving a data request, challenge a requester to provide all in a series of secure tokens. A requester would be instructed to visit a specific server in order to retrieve the needed token, that server being remote and discrete from the server being protected. Upon visiting that server, the requester would receive the first of many required tokens in order to be authorized to be served with content. The server would provide a token as well as the address of the next server in the series which must be contacted to find the next token along with instructions on how to find the next token, somewhere on the Internet. The requester is sent on a type of "quest" in which a series of tokens must be retrieved from 10-20 servers located in various countries in which the only way to find all of the tokens is to visit all of the servers. The tokens are regularly rotated and may vary depending upon the IP of the requester. This ensures that latency is amplified for all of the nodes participating in the DDoS attack and also ensures that a single, successfully completed "token quest's" results cannot be used by an attacker to allow all nodes authorization to make malicious data requests from the target server.

Conclusion

While it might ordinarily 50ms for a webpage to begin to load, when receiving this type of token challenge, it might take 5000ms to accomplish the same task. For a legitimate user, this would not present such an impediment so as to prevent normal use of a system. Because most attackers are not willing to wait 5000ms before reiterating requests, these requests can be identified and blocked in real-time. If an attacker were to be patient enough to wait five seconds to receive a single block of data such as a webpage's contents, they would be unlikely to be able to sap a sufficient amount of bandwidth from the primary server in order to impact performance.